

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application for:

**DYNAMIC COMPOSITION OF PRE-ENCRYPTED
VIDEO ON DEMAND CONTENT**

Inventor(s): Leo Mark Pedlow, Jr. and Davender Agnihotri

Docket Number: SNY-T5712.02

Prepared By: Miller Patent Services
2500 Dockery Lane
Raleigh, NC 27606

Phone: (919) 816-9981
Fax: (919) 816-9982
Email: miller@patent-inventions.com

CERTIFICATE OF EXPRESS MAILING FOR NEW PATENT APPLICATION

"Express Mail" mailing label number: ER999163695US

Date of Deposit: 3/16/04

I Hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450.

Typed or printed name of person mailing paper or fee: Catherine N. Miller

Signature of person mailing paper or fee: Catherine N. Miller

5

DYNAMIC COMPOSITION OF PRE-ENCRYPTED VIDEO ON DEMAND CONTENT

10

CROSS REFERENCE TO RELATED DOCUMENTS

This application is related to and claims priority benefit of U.S. Provisional Patent Application Serial No. 60/516,050 filed October 31, 2003 to Pedlow et al. for "Dynamic Composition of Pre-Encrypted Video On Demand Content" which is hereby incorporated
15 by reference. This application is also related to U.S. Patent Applications docket number SNY-R4646.01 entitled "Critical Packet Partial Encryption" to Unger et al., serial number 10/038,217; patent applications docket number SNY-R4646.02 entitled "Time Division Partial Encryption" to Candelore et al., serial number 10/038,032; docket number SNY-R4646.03 entitled "Elementary Stream Partial Encryption" to Candelore,
20 serial number 10/037,914; docket number SNY-R4646.04 entitled "Partial Encryption and PID Mapping" to Unger et al., serial number 10/037,499; and docket number SNY-R4646.05 entitled "Decoding and Decrypting of Partially Encrypted Information" to Unger et al., serial number 10/037,498 all of which were filed on January 2, 2002 and are hereby incorporated by reference herein.

25

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent document or the patent disclosure, as it appears in the Patent

and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND

5 The Passage™ initiative (Passage is a trademark of Sony Electronics Inc.), promoted by Sony, provides a mechanism for MSOs (Multiple Service Operators) to deploy non-legacy headend equipment, subscriber devices and services on their existing legacy networks. At present, in the USA, these networks are most commonly supplied by either Motorola (formerly General Instrument) or Scientific Atlanta. These two
10 companies at present constitute better than a 99% share of the US cable system market as turnkey system providers. The systems, by design, employ proprietary technology and interfaces precluding the introduction of non-incumbent equipment into the network. An MSO, once choosing one of these suppliers during conversion from an analog cable system to a digital cable system, faces a virtual monopoly when seeking suppliers for
15 additional equipment as their subscriber base or service offering grows.

 Before the Passage™ initiative, the only exit from this situation was to forfeit the considerable capital investment already made with the incumbent provider, due to the intentional incompatibility of equipment between the incumbent and other sources. One primary barrier to interoperability is in the area of conditional access (CA) systems, the
20 heart of addressable subscriber management and revenue collection resources in a modern digital cable network.

 The Passage™ technologies were developed to allow the independent coexistence of two or more conditional access systems on a single, common plant. Unlike other attempts to address the issue, the two systems operate with a common transport stream
25 without any direct or indirect interaction between the conditional access systems. Some of the basic processes used in these technologies are discussed in detail in the above-referenced pending patent applications.

 The above-referenced commonly owned patent applications, and others, describe inventions relating to various aspects of methods generally referred to herein as partial

encryption or selective encryption, consistent with certain aspects of Passage™. More particularly, systems are described therein wherein selected portions of a particular selection of digital content are encrypted—using two (or more) encryption techniques while other portions of the content are left unencrypted. By properly selecting the portions to be encrypted, the content can effectively be encrypted for use under multiple decryption systems without the necessity of encryption of the entire selection of content. In some embodiments, only a few percent of data overhead is consumed to effectively encrypt the content using multiple encryption systems. This results in a cable or satellite system being able to utilize Set-top boxes (STB) or other implementations of conditional access (CA) receivers (subscriber terminals) from multiple manufacturers in a single system - thus freeing the cable or satellite company to competitively shop for providers of Set-top boxes.

In each of these disclosures, the clear content is identified using a primary Packet Identifier (PID). A secondary PID (or shadow PID) is also assigned to the program content. Selected portions of the content are encrypted under two (or more) encryption systems and the encrypted content transmitted using both the primary and secondary PIDs (one PID or set of PIDs for each encryption system). The so-called legacy STBs operate in a normal manner decrypting encrypted packets arriving under the primary PID and ignoring secondary PIDs. The newer (non-legacy) STBs operate by associating both the primary and secondary PIDs with a single program. Packets with a primary PID are decoded normally and packets with a secondary PID are first decrypted then decoded. The packets associated with both PIDs are then assembled together to make up a single program stream. The PID values associated with the packets are generally remapped to a single PID value for decoding (e.g., shadow PIDs remapped to the primary PID value or vice versa.)

BRIEF DESCRIPTION OF THE DRAWINGS

Certain illustrative embodiments illustrating organization and method of operation, together with objects and advantages may be best understood by reference

detailed description that follows taken in conjunction with the accompanying drawings in which:

FIGURE 1 is a block diagram of a clear video VOD system.

FIGURE 2 is a diagram illustrating storage of I-frame data to support trick mode
5 operation in a VOD system.

FIGURE 3 is a block diagram of a pre-encrypted VOD system using a single (legacy) encryption system.

FIGURE 4 is a block diagram depicting a hybrid composite VOD system architecture consistent with certain embodiments of the present invention.

10 **FIGURE 5** is a block diagram of a re-encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 6 illustrates a dynamic composition pre-encrypted VOD architecture consistent with certain embodiments of the present invention.

15 **FIGURE 7** also illustrates a dynamic composition pre-encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 8 illustrates a dynamic composition pre-encrypted VOD architecture using dual trick play indices consistent with certain embodiments of the present invention.

20 **FIGURE 9** is a flow chart of a first exemplary process consistent with certain embodiments of the present invention.

FIGURE 10 is a flow chart of a second exemplary process consistent with certain embodiments of the present invention.

25 **ACRONYMS, ABBREVIATIONS AND DEFINITIONS**

ASI - Asynchronous Serial Interface

CA - Conditional Access

CASID - Conditional Access System Identifier

CPE - Customer Premises Equipment

DHEI - Digital Headend Extended Interface

ECM - Entitlement Control Message

EPG - Electronic Program Guide

GOP - Group of Pictures (MPEG)

5 **MPEG** - Moving Pictures Experts Group

MSO - Multiple System Operator

PAT - Program Allocation Table

PID - Packet Identifier

PMT - Program Map Table

10 **PSI** - Program Specific Information

QAM - Quadrature Amplitude Modulation

RAID - Redundant Array of Independent Disks

RAM - Random Access Memory

SAN - Storage Area Network

15 **VOD** - Video on Demand

Critical Packet - A packet or group of packets that, when encrypted, renders a portion of a video image difficult or impossible to view if not properly decrypted, or which renders a portion of audio difficult or impossible to hear if not properly decrypted. The term “critical” should not be interpreted as an absolute term, in that it may be possible to hack
20 an elementary stream to overcome encryption of a “critical packet”, but when subjected to normal decoding, the inability to fully or properly decode such a “critical packet” would inhibit normal viewing or listening of the program content.

Selective Encryption (or Partial Encryption) – encryption of only a portion of an elementary stream in order to render the stream difficult or impossible to use (i.e., view
25 or hear).

Dual Selective Encryption – encryption of portions of a single selection of content under two separate encryption systems.

Passage™ - Trademark of Sony Electronics Inc. for various single and multiple selective encryption systems, devices and processes.

Trick mode – an operational mode of playback of digital content to simulate fast forward, rewind, pause, suspend (stop), slow motion, etc. operations as in a video tape system.

The terms “a” or “an”, as used herein, are defined as one, or more than one. The
5 term “plurality”, as used herein, is defined as two or more than two. The term “another”,
as used herein, is defined as at least a second or more. The terms “including” and/or
“having”, as used herein, are defined as comprising (i.e., open language). The term
“coupled”, as used herein, is defined as connected, although not necessarily directly, and
not necessarily mechanically. The term “program”, as used herein, is defined as a
10 sequence of instructions designed for execution on a computer system. A “program”, or
“computer program”, may include a subroutine, a function, a procedure, an object
method, an object implementation, in an executable application, an applet, a servlet, a
source code, an object code, a shared library / dynamic load library and/or other sequence
of instructions designed for execution on a computer system.

15 The terms “scramble” and “encrypt” and variations thereof may be used
synonymously herein. Also, the term “television program” and similar terms can be
interpreted in the normal conversational sense, as well as a meaning wherein the term
means any segment of A/V content that can be displayed on a television set or similar
monitor device. The term “storing” as used herein means both the act of placing data into
20 a storage medium and holding the data in storage in the storage medium. The term
“video” is often used herein to embrace not only true visual information, but also in the
conversational sense (e.g., “video tape recorder”) to embrace not only video signals but
associated audio and data. The term “legacy” as used herein refers to existing technology
used for existing cable and satellite systems. The exemplary embodiments of VOD
25 disclosed herein can be decoded by a television Set-Top Box (STB), but it is
contemplated that such technology will soon be incorporated within television receivers
of all types whether housed in a separate enclosure alone or in conjunction with recording
and/or playback equipment or Conditional Access (CA) decryption module or within a
television set itself.

DETAILED DESCRIPTION

While this invention is susceptible of embodiment in many different forms, there is shown in the drawings and will herein be described in detail specific embodiments, with the understanding that the present disclosure of such embodiments is to be considered as an example of the principles and not intended to limit the invention to the specific embodiments shown and described. In the description below, like reference numerals are used to describe the same, similar or corresponding parts in the several views of the drawings.

CLEAR VOD ARCHITECTURES

The decision on a particular VOD architecture is the result of the interaction between a complex set of both independent and dependent variables, providing a solution to an equation of state. Some of the variables are fixed directly as a result of choices by the MSO. Others are constrained by factors such as the existing incumbent system, location, size, available capital and ROI requirements.

A generalized VOD system 10, as shown in **FIGURE 1**, contains some or all of the following elements / resources: Content Aggregation and Asset management 14, Content distribution (SAN) 18, Video server module(s) 22, Session Management 26, Transaction management 30, Billing system 34, EPG server or VOD catalog server 38, Transport router/switch fabric (routing matrix) 42, Stream encryption device(s) (not shown in this Figure), and QAM modulators/upconverters and other edge resources 46. This VOD system 10 provides programming to the subscriber terminals such as 50 for ultimate viewing and listening on a TV set or other monitor device 54.

In operation, content is received from various sources including, but not limited to, satellite broadcasts received via one or more satellite dishes 58. Content is aggregated at 14 and cataloged at EPG server or VOD catalog server 38. Content is then distributed at 18 to one or more video servers 22. When a subscriber orders a VOD selection, a message is sent from the subscriber terminal (e.g., STB) 50 to the session manager 26.

The session manager 26 notifies the transaction manager 30 to assure that the billing system 34 is properly brought into play. The session manager 26 selects a VOD server from a cluster of VOD servers having the requested content on it and having a signal path that reaches the node serving the subscriber. The session manager also enables the
5 routing matrix 42 to properly route the selected video content through the correct edge resources 46 for delivery to the subscriber terminal 50.

TRICK MODES

One aspect of VOD that has become a “signature” feature is the support of “trick
10 modes”. These are operational modes invoked by the session client that mimic a traditional VCR or DVD player and includes fast forward, rewind, pause, suspend (stop), slow motion, etc. Trick modes have been heretofore implemented through the creation of multiple files containing a subset of the original content (subfiles) as illustrated in **FIGURE 2**. The content is generally stored in a set of RAID drives 70. A particular
15 selection of content is stored in its entirety in a file 74 within the RAID drives 70. A set of subfiles for rewind and fast forward trick modes (files 78 and 80 respectively) contain I-frames ordered in a manner that will permit playback sequentially to achieve the rewind and fast forward effect. Typically, these subfiles contain only I-frames, since I-frames contain stand-alone whole pictures (see ISO/IEC 13818-2, section 6.1.1.7). I-frames are
20 somewhat larger than B or P frames, and they typically represent approximately as much as 21% of the data in a given video selection.

A file containing only I-frames extracted from the original content affords the ability to have accelerated playback, since typical GOP (group of pictures) structures have only one frame in about 10 to 20 as an I-frame. If the I-frame files are played at
25 normal rates (1 frame per 33 mS) the pictures will appear to the viewer to sequence at about a 10x to 20x rate, though the actual data rate is the same as the original content. If the I-frame sequence is reversed in the file, the motion will appear to run backwards. This is the method used to implement fast forward and rewind trick modes.

By attaching an index count to match the I-frames in the original content file to the duplicated I-frames stored in the associated subfiles 78 and 80, a method is provided to allow immediate transition from normal speed forward play to fast forward or rewind. In operation the video server plays the selected content file and upon subscriber selection
5 of a trick mode (or vice versa) the server notes the index value of the closest I-frame and then opens the appropriate associated subfile 78 or 80 and moves to the I-frame in the subfile with the same corresponding index. The video server treats all stream content (main file or subfiles) the same and always spools the MPEG packets to the outgoing transport stream at the same constant bit rate through multiplexers and buffers 84 as
10 shown. It is through this method that trick modes are typically implemented on a slotted, session based system without the encumbrance of additional, dynamic bit rate issues.

Unfortunately, the use of such multiple subfiles results in storage space inefficiencies. As will be seen, these inefficiencies can become compounded in systems utilizing multiple encryption.

15

VOD PROGRAM SPECIFIC INFORMATION

A function of the VOD video server(s) 22, in addition to origination of session A/V content, is the creation of the associated, session specific PSI (program specific information). This information is a departure from the broadcast model in that the PSI is
20 extremely dynamic. The content of the PAT and subordinate PMTs change whenever a new session is started or ended. In the broadcast world, the PSI changes very seldom because the PSI tables reflect only the structure of the transport multiplex, not the actual A/V content carried within.

The VOD video server 22 dynamically assigns a new session to an existing,
25 available "slot" in an outgoing transport multiplexed stream. The slot is denoted by the MPEG program number and in many cases, the combination of which transport stream (TSID) and program number determine at the service level a unique session and the routing that occurs as a result. Edge resources 46 generally are not configured dynamically. The routing of content appearing on a particular input port to a specific

QAM carrier at the output is determined through a preconfigured, static assignment of TSID/input port and program number mapping to specific QAM resources in the device. This same mapping information is also loaded in the VOD system so that once a session is requested by and authorized for a specific subscriber terminal 50, a solution to a routing matrix 42 can be determined to find the appropriate VOD server 22 and QAM transport 46 serving the requestor. This solution also considers dynamic issues such as which servers 22 the requested asset is loaded upon, and server loading/available slots in addition to the simpler, static solution to finding the first possible path to the requesting subscriber terminal 50.

10 In addition to solving the routing matrix 42 and provisioning the session with PIDs and PSI appropriate to follow the intended route, elements of the same information (program ID and QAM frequency) are also communicated to the session client at subscriber terminal 50 at the subscriber's premises so that the requested stream can be properly received and presented to the subscriber.

15

CLEAR VOD DISTRIBUTION

Perhaps the simplest VOD distribution system implementation is a clear VOD distribution system, i.e. one that contains no encryption as depicted in **FIGURE 1**. While not providing any safekeeping of what might be considered the entertainment medium's most valuable properties, namely current feature films, etc., clear VOD avoids many of the issues that the incumbent cable system providers to date have not adequately addressed and that introduction of a second, alternative CA system complicates even further still. Various arrangements for providing selective or full encryption in a VOD environment are discussed below. Throughout this discussion, it is instructive to carry an example VOD movie through the various embodiments to illustrate the relative storage efficiencies obtained with the various systems disclosed. A real world example of a VOD movie which will be used throughout this document has the following attributes:

Compressed video data rate: 3Mbit/S

Movie length: 120 minutes (2 Hrs)

I-frame overhead: 17%
Total storage used for
the video portion of a
single, clear (unencrypted)
5 copy of a film: 3.618GBytes.

PRE-ENCRYPTED VOD DISTRIBUTION

Pre-encrypted VOD systems such as system 100 shown in **FIGURE 3** can be architecturally similar to clear VOD distribution systems. One difference between the
10 two is that on pre-encrypted systems there is pre-processing of the content prior to storage in the VOD system to provide safekeeping of content during the storage and distribution phases. This pre-processing can be carried out in pre-encryptor 104. Data security is implemented through storage of previously encrypted content within the video server(s) 22. While the clear VOD system contains directly viewable MPEG or other
15 compressed A/V content on the server(s) 22, the pre-encrypted model stores this same content in a form that is only decipherable using a properly entitled subscriber terminal 50.

The pre-encryption process can be performed by the MSO at the time of deployment on the VOD system 100, prior to loading into the storage area network
20 (SAN) used to propagate content to all of the video servers in the MSO's system. Alternatively, the encryption may be performed prior to receipt of the content by the MSO at an external service bureau, content aggregator or by the distributor or studio. In this case, the content is theoretically secured throughout the distribution phase, storage phase and transmission to subscriber for display on an authorized device. The use of pre-
25 encryption prior to distribution of content to the MSO potentially adds to the complexity of entitlement distribution, separate from the content distribution, for installation on the VOD transaction manager 30 to allow bone fide subscribers to decrypt the purchased content. For purposes of this document, content will be considered stored in the VOD

video server if it is stored either directly in the VOD video server or indirectly in the VOD video server (i.e., is accessible by the VOD video server).

Many pre-encrypted VOD architectures share one or more of the following common drawbacks:

- 5 • Additional handling of new content may be needed to perform the pre-encryption prior to loading into the server, either by the MSO or service bureau.
- Coordination and/or distribution is required for entitlements matching the access criteria used to encrypt the content stored in the server.
- Limited “shelf life” of the encryption keys used to secure the stored content,
10 rendering decryption impossible at a later date.
- Incapability of present VOD video servers to load pre-encrypted streams.
- Incompatibility of pre-encrypted streams with present methods supporting trick mode play (fast-forward & rewind) on screen.
- One common key is used for all sessions accessing a particular program and it
15 remains the same for the duration of time the content is in inventory on the server.
- According to MSOs familiar with the subject, pre-encrypted VOD streams are unsupported by conditional access technologies from certain manufacturer(s).

The issue regarding trick play and pre-encryption is based upon the concept that
20 VOD servers 22 currently expect clear content and then subsequently identify the I-frames and store or otherwise segregate them for access in fast-forward or fast rewind playback modes, as described in conjunction with **FIGURE 2**. If the stream is pre-encrypted prior to storage upon the server, it may be difficult or impossible for the server
22 to examine packet payloads to identify I-frames during the process of importation into
25 the server 22 to create trick mode files 78 and 80 or associated indices. Many current systems will not accept streams for importation that are pre-encrypted.

SEGREGATED STORAGE PRE-ENCRYPTION

A segregated storage mechanism can be physically similar to the architecture of the clear VOD distribution system. The content is encrypted in its entirety (100%) and a separate copy of the complete feature is stored for each different conditional access
5 format supported by the MSO. The organization and configuration of the system is such that when a subscriber initiates a session on the server, the stream files for the selected content containing the CA format appropriate to the specific equipment deployed at the subscriber's premises requesting the session are spooled and delivered. This method offers a low system complexity encrypted VOD system but may suffer from some of the
10 same issues common to other pre-encryption topologies, mentioned previously. In addition, a very significant storage penalty (one or more encrypted duplicate copies of the same movie) is incurred.

If one refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require an additional
15 7.236GBytes to store using segregated pre-encryption supporting two different CA systems.

Changes to the method employed by the VOD system are used for creating dynamic PSI data to implement this architecture supporting multiple CA systems. The VOD system session manager is made aware of which conditional access method is
20 appropriate for a session requested by a specific subscriber. This information is in turn transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server is cognizant of the conditional access resources (ECMs) for each program stored on the server and these resources can be dynamically allocated on unique
25 PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

COMPOSITE STORAGE PRE-ENCRYPTION

Composite storage is essentially the storage on the video server of a selectively encrypted stream such as a Passage™ processed stream that contains previously encrypted “critical packets” for a plurality (two or more) of independent conditional access systems (i.e., dual selective encrypted). The stream may be prepared identically to the processing of a selectively encrypted broadcast stream as described in the above-referenced pending patent applications, except that the resultant transport stream is recorded to a hard disk or other suitable computer readable storage medium, instead of being sent directly to a QAM modulator for HFC distribution to the requesting subscriber. As with other pre-encryption models, the content can be encrypted by either the MSO at time of deployment on the VOD system, a third party service bureau, by the studios themselves (the latter two cases being prior to receipt of the content by the MSO), or by or under control of other entities.

In this embodiment the small additional overhead in content storage (typically 2% – 10% representing “critical packets” that are multiple encrypted) is traded for the support of multiple independent CA formats without replication of entire streams. A negative aspect, in addition to those mentioned previously and common to other pre-encryption topologies, is the vulnerability of the prepared selectively encrypted stream to corruption by downstream equipment containing transport remultiplexing functionality that is not specifically designed to maintain the integrity of the selective encryption process applied to the stream.

If one refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require approximately 3.690GBytes to store using composite storage pre-encryption supporting two different CA systems with a critical packet “density” of 2%.

Certain changes to the method employed by the VOD system for creating dynamic PSI data can be used to implement this architecture. The VOD system session manager can be made to be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the

video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server is cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, can indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the video server dynamically allocates another set of PIDs for the shadow packets associated with the respective audio and video component streams for each session in the manner described in the above-referenced patent applications. This information can be included in the PSI sent in sessions requested by non-legacy clients. In total, eight different PIDs and corresponding data resources are dynamically allocated and managed by the server for each session: PAT (one table common to all sessions, but modified for each), PMT, Primary Video, Primary Audio, Shadow Video, Shadow Audio, Legacy ECM and Alternative ECM. Six of these entities can be stored in the embedded stream and use dynamic PID remapping for each session.

Consider the issue of which device to use in conjunction with performing the legacy encryption of the "critical" packets prior to storage on the VOD video server. If the legacy device is specially designed to process content destined for loading into a VOD video server, it may not accept a selectively encrypted stream at its input. The content format specified for VOD servers often uses a single program transport multiplex containing a single PAT entry, single PMT entry and service components, for one audio and one video stream. The shadow packets added in a composite selectively encrypted transport stream may prove problematic for a legacy VOD pre-encryption device, in certain instances. It is more probable that a device or process (since there are no real time requirements, an off-line process running on a PC or UNIX server may suffice) to process a candidate stream before passing through the legacy pre-encryptor and then post-encryption reconcile to extract only the encrypted "critical" packets for insertion

into the VOD video server 22. The same or similar algorithms and techniques for performing this manipulation for selective encryption processing as described in the above-referenced patent applications can be adapted to VOD applications for off-line work.

5 The VOD server 22 may also be modified to allow introduction of streams having multiple service elements (primary video, primary audio, shadow video, shadow audio) uniquely associated with a Passage™ transport. The present video servers generally only allow one each, primary video and audio, respectively. The quartet of data representing Passage™ processed A/V content should preferably be managed as a indivisible set on
10 the VOD video server 22.

 Some additional bandwidth efficiencies may be obtained if, at the edge resources, shadow packets are removed from the composite streams in sessions serving legacy clients. Similarly, in certain embodiments, the edge resources, if selective encryption aware, could reinsert the shadow packets embedded in the stored stream in place of the
15 legacy encrypted packets on the original program PID. These improvements would result in no carriage overhead for support of multiple conditional access systems on a single transport.

HYBRID COMPOSITE STORAGE PRE-ENCRYPTION

20 Hybrid composite storage is a variant of the composite storage concept, but incorporates elements of session-based encryption for implementing the alternative conditional access encryption. In this scenario, depicted as system 130 of **FIGURE 4**, the legacy “critical” packets, comprising approximately 2-10% of the total content, are pre-encrypted by the legacy conditional access system 104 using selective encryption
25 technology for managing the process. The selective encryption is managed in selective encryption processor 134. The duplicate copy of “critical” packets, which are located on previously unused PIDs, is left unencrypted. This latter aspect is the departure from the composite storage scenario described above. The composite stream of unencrypted non-critical packets, legacy encrypted “critical” packets on the original service PIDs and an

unencrypted, duplicate copy of the “critical” packets on alternate service PIDs is stored on the video server 22 as a single stream.

Upon playback to a subscriber session, if the session is destined for a legacy STB (represented by subscriber terminal 50), the existing paradigm for pre-encrypted content is followed and no special action is taken. The stream is routed at routing matrix 138 operating under control of session manager 26, through a session encryption device 142 capable of performing encryption using the alternative conditional access system 144, but the session manager 26 does not provision the device to perform encryption on elements of the stream and it is sent directly to the requesting subscriber without further modification. To maintain security of the outgoing stream and to reduce the bandwidth of the session for legacy sessions, the stream is processed through an add-drop remultiplexer 148 and the clear “critical” content on alternate service PIDs are removed from the outgoing transport. The output stream is then routed at routing matrix 152 to appropriate edge resources 46 for delivery to the subscriber terminal 50. In one embodiment, the session encryption device 142 that performs encryption using the alternative conditional access system also contains the add-drop multiplexer capability. Other variations will also occur to those skilled in the art upon consideration of the present teaching.

If, on the other hand, the session is destined for a non-legacy STB (also as represented in this illustration by subscriber terminal 50), the stream is routed through session encryption device 142 capable of performing encryption using the alternative conditional access system and only the “critical” packets on alternate service PIDs (previously in the clear) are encrypted using the alternative conditional access system 144, as provisioned by the session manager.

Some additional bandwidth efficiencies may be obtained for these non-legacy sessions, if the edge device is selective encryption aware, by reinserting the shadow packets embedded in the stored stream, now encrypted, in place of the legacy encrypted packets on the original program PID. This improvement would result in no carriage overhead for support of multiple conditional access systems on a single transport.

In certain embodiments, a preprocessor can be used to perform selective encryption of content to be loaded onto the video server. A modified file protocol can be used to allow the video server to import and associate these files. Either the preprocessor or the video server can be designed to perform the indexing. An alternate instantiation
5 could be use to perform all selective encryption pre-processing (e.g., PID mapping and packet duplication) within the VOD video server 22 itself. This could be accomplished by modifying the VOD video server 22 application to add a pre-processor task as a separate executable, called by the VOD video server 22 during the process to prepare content for pre-encryption.

10 Changes can be implemented to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager 26 is made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information can in turn be transferred to the VOD video server 22 that has been selected as the source for the session so that the
15 appropriate PSI can be created for the session, including conditional access specific data. The VOD video server 22 is cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, can indicate the
20 appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the VOD video server 22 dynamically allocates PIDs for the shadow packets associated with the respective audio and video component streams for each session. This information is included in the PSI sent in sessions requested by non-legacy
25 clients. Just like in the more general composite storage architecture discussed in the previous section, the video server manages multiple resources and PIDs. The hybrid topology reduces the unique entities by one from eight to seven: there is no need for alternative ECM PID or data resource in the stored composite stream. This information

will be added later in a downstream device providing the alternative conditional access encryption for those sessions destined for decoding upon a non-legacy client.

RE-ENCRYPTED DISTRIBUTION

5 A hybrid approach is provided in a re-encrypted distribution architecture. This topology leverages the paradigms established for pre-encrypted content preparation, storage, management, etc. but adds support for session based encryption for the alternative conditional access systems added to an existing incumbent system. Referring to the exemplary embodiment of **FIGURE 5**, a legacy decryption device 182, operating
10 to decrypt using the legacy CA system 184, is added to the transport stream path exiting the VOD video server 22 (via routing matrix 186). After the decryption device 182, the transport stream passes through a contemporary session based encryption device 188 based upon the alternate CA system. The VOD session manager 26, on a session-by-session basis, determines which sessions will pass through the decryption device 182
15 intact and be modulated and transmitted to the subscriber unaltered. A path 190 between the routing matrices preserves the pre-encrypted content and delivers it to subscribers having legacy equipment. In either case, the output stream passes through routing matrix 152 to the appropriate edge resources for delivery to the subscriber terminal 50.

 Alternatively, the VOD system session manager 26, through interaction with both
20 legacy CA system 184 and alternate CA system 194, can both actuate the decryption device 182 and activate session based encryption device 188 for a particular session, thereby supporting subscribers with non-legacy equipment at their premises. Thus, this system 180 can support either legacy or non-legacy (alternate CA) encryption.

 Certain embodiments of this architecture support pre-encryption on legacy
25 systems not presently supporting session-based encryption, while providing the ability to deliver session based encryption for the alternative CA system 194 integrated into the existing legacy network. Certain embodiments of this architecture may face some of the same issues as mentioned previously and common to other pre-encryption topologies. In addition, it experiences the additional cost burden of a legacy decryption element and the

challenges of dynamically configuring and operating such a device. There may be additional costs faced in a specific deployment for switching and routing equipment that may be necessary to move transport streams “around” the legacy decryption device. However, this architecture permits storage of fully encrypted content to safeguard the content while enabling dual encryption without storage penalty.

Changes can be made to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager can be made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server can be made to be cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

In this example, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using re-encryption supporting two different CA systems.

Thus, in certain embodiments consistent with the present invention, a method of storage and distribution of video-on-demand content, involves receiving a request from a subscriber terminal to transfer the selection of video content to the subscriber terminal; determining that the subscriber terminal is able to decrypt content encrypted under the first encryption system or under a second encryption system; if the subscriber terminal is able to decrypt the content encrypted under the first encryption system, then routing a selection of content that has been encrypted under the first encryption system to the subscriber terminal; if the subscriber terminal is able to decrypt the content encrypted under the second encryption system, then: a) decrypting the selection of content

encrypted under the first encryption system to produce clear content; b) encrypting the selection of content under the second encryption system to produce a re-encrypted selection of content; and c) routing the re-encrypted selection of content to the subscriber terminal.

5 In other words, a method of storage and distribution of video-on-demand content consistent with certain embodiments involves receiving a request from a subscriber terminal 50 to transfer the selection of video content to the subscriber terminal 50. If the subscriber terminal is able to decrypt the content encrypted under the first encryption system, the encrypted content is routed to the subscriber terminal 50. If the subscriber
10 terminal is able to decrypt the content encrypted under the second encryption system, the content is first decrypted and then re-encrypted under the second encryption system before routing to the subscriber terminal 50.

DYNAMIC COMPOSITION PRE-ENCRYPTION

15 Another pre-encrypted VOD architecture consistent with certain embodiments of the present invention is dynamic composition pre-encryption. In this scheme, each program or movie is stored in three or more elements on the VOD video server 22. Referring to **FIGURE 6**, clear content is stored at 200. Critical packets are selected according to a suitable selection criterion associated with the selective encryption
20 process. Thus, the content that is stored has either “critical” packets or non-critical packets. The “critical” packets generally constitute approximately 2% to 10% of the program (depending upon program content and the selection criteria used to select packets for encryption) and are encrypted. A separate copy of the critical content is maintained for each conditional access system supported by the MSO. In this illustration,
25 for example, the critical packets associated with a first CA system (CA1) is stored at 202 while encrypted content associated with CA 2 is stored at 206. By using a selection criterion that involves selection of certain I-frames, the fast forward I-Frames can be made to incorporate the encrypted content and stored together as encrypted I-frames 210 (and 206). The packets in both the “critical” packet fast forward file 210 as well as the

clear (unencrypted), non-critical packet file 200 are indexed to maintain temporal correlation between the two files. These indices either may be monotonic packet counts from start of stream or calculated packet offsets from the last PCR.

When a subscriber session is initiated, the main file 200 containing the clear content, less "critical" and fast forward packets, is queued in the video server for playout. In addition, the file containing the "critical" and fast forward packets 210, pre-encrypted in the CA format appropriate for the CPE of the subscriber requesting the session, is also queued for playout. When the program playback is started, the video server reconstructs a single program multiplex in its stream buffer feeding the outgoing transport the correct sequence of packets based upon the indices in the two component files. Although, in general, only about 2-10% of the packets are encrypted in a selective encryption system according to the above pending patent applications, even further security is provided by encryption of all of the I frames in the present embodiment. Rewind I-frames can be stored either as encrypted or unencrypted packets. A dual selective encrypted embodiment is depicted in **FIGURE 7**.

While the external composition and data flow appears similar to the clear VOD system depicted in **FIGURE 1**, the internal architecture of the video server changes significantly, as shown in the exemplary storage architecture of **FIGURE 6** and **FIGURE 7**.

Certain embodiments of this method offer several distinct advantages that may not be readily apparent. The stream files containing "critical" packets may be the same one as the extracted subfile containing all I-frames for "trick" modes, as was described previously in the general discussion of VOD system architecture. If this opportunity is taken, then a storage economy can be realized over all pre-encrypted schemes including traditional (unencrypted) VOD, as deployed today. The traditional VOD video server has three files for each feature or movie: two containing just I-frames (one in reverse order) and one containing the complete original copy. Research on encoded streams conducted by Sony has shown that the I-frames typically represent between 12%-21% of the total content, typically around 17%. With the dynamic composition method, if the "critical"

packet files are chosen to contain complete I-frames, a separate file of critical data used solely for encryption purposes is no longer necessary, saving 2% to 10% storage for this method. In addition, since this method removes the redundant I-frames from the clear stream file, an additional (nominal) 17% storage savings is also realized. This indicates a potential 27% nominal (31% maximum) video server disk storage savings for a single CA system model over the composite storage model VOD system described above.

When compared to the segregated storage model described above, one entire duplicate copy of a program can be eliminated and the addition of one additional CA format adds no storage or bandwidth overhead when compared to a traditional clear VOD server implementation. The reason for the “free” second CA format is that the 17% nominal storage saving realized by using the same I-frame file for both fast forward “trick” modes and “critical” content used for selective encryption is consumed by replicating just the I-frame file and encrypting it with the alternative CA format.

15 DYNAMIC COMPOSITION PRE-ENCRYPTION WITH FORWARD AND REVERSE INDEXING

If one takes the concept of dynamic composition pre-encryption described above one step further, the current convention in VOD systems to store the same I-frames of a movie in forward and reversed sequence to allow fast forward and rewind “trick” modes can be eliminated. An illustration of this concept is shown in the example of **FIGURE 8**. These dual files for forward and reverse are replaced by a single file 320 of I-frames in normal forward sequence with two sets of indices, one set 322 for playing the I-frame file in forward order and one set 324 for playing the I-frame file in reverse order. The appropriate sets of indices are chosen depending on whether forward or reverse high-speed motion is desired. The forward indices are also used to reconstruct the normal speed stream when matching the I-frame file to the non-critical content file to reconstruct the entire stream. On a clear or re-encrypted VOD system, this will allow up to about 21% storage savings. On a composite pre-encrypted storage system, up to about 42% storage savings may be realized.

It is noted that although the arrangement of **FIGURE 7** provides substantial savings in storage space over the techniques currently in use, additional savings in storage space can be realized by the recognition that the information stored in the trick mode content file is redundant to the I frames stored in the normal play content file 200.

5 By spooling normal play content from both files, an additional savings of up to approximately 21% can be realized as depicted in **FIGURE 8**. In this illustration, all I-frame data (intra-coded data) are stored in the trick mode content file 320, and supplemental normal play content (inter-coded data, B and P frame data) is stored in the normal play content file 300. The bidirectional indices concept is extended for even

10 further storage economy in this embodiment. If one recognizes that the normal mode playback file contains a duplication of the same I-frames played in "trick" modes, a dynamic architecture can be created to remove any redundant I-frame content from the normal mode playback file. During normal playback, the two files are "blended" (normal play and "trick" modes), while only the I-frame sequences in the "trick" mode file are

15 accessed during fast forward, fast reverse (rewind), etc.

Thus, according to certain embodiments consistent with the present invention, a method of processing digital video content, wherein the digital video content comprises intra-coded frames and inter-coded frames, involves selecting a plurality of the intra-coded frames for encryption to produce selected frames; encrypting the selected frames

20 under a first encryption algorithm to produce first encrypted frames; storing the inter-coded frames in a first file; and storing the intra-coded frames, whether encrypted under the first encryption algorithm or unencrypted, in a second file. For a multiple encryption embodiment consistent with the present invention, the method further involves duplicating the intra-coded frames; encrypting duplicates of the selected frames under a

25 second encryption algorithm to produce second encrypted frames; storing the intra-coded frames, whether encrypted under the second encryption algorithm or unencrypted, in a third file.

As noted above, the current convention in VOD systems to store the same I-frames of a movie in forward and reversed sequence to allow fast forward and rewind

“trick” modes can be eliminated. An illustration of this concept is shown in the example of **FIGURE 8**. These dual files for forward and reverse are replaced by a single file 320 of I-frames in normal forward sequence with two sets of indices, one set 322 for playing the I-frame file in forward order and one set 324 for playing the I-frame file in reverse order, or equivalently, by a single index that is traversed in the forward or reverse direction for FF or FR play respectively. The appropriate sets of indices are chosen depending on whether forward or reverse high-speed motion is desired. The forward indices are also used to reconstruct the normal speed stream when matching the I-frame file to the non-critical content file to reconstruct the entire stream. On a clear or re-encrypted VOD system, this will allow up to about 21% storage savings. On a composite pre-encrypted storage system, up to about 42% storage savings may be realized.

If the “trick” mode subfile and the “critical” data encrypted content file can be the same, the content is selectively encrypted up to approximately a nominal 17% level (~21% max), much higher than the commonly proposed Passage™ encryption level of approximately 2%, but carrying no inherent storage or system capacity costs, as do other schemes. It should be noted that all or only selected ones of the I frames can be encrypted according to certain embodiments consistent with the present invention. For this system to work, some changes to the video server software design might be necessary, but these changes would be modifications to the existing processes and would not require substantial new development on the part of the server vendor.

In the example of **FIGURE 8**, either two reference tables or one could be used in implementing various embodiments consistent with this example. In this example, however, it should be remembered that the normal play file does not contain a full set of content, but rather may contain only data associated with intra-coded frames. Thus, to carry out a normal playback, the index tables are used to identify a full set of data and data are pulled from both file 300 and file 320.

In any case, the indices can be visualized as a table such as **TABLE 1** and **TABLE 2** below. In the case of the using a forward and a reverse index, **TABLE 1** represents the forward index and **TABLE 2** represents the reverse index.

TABLE 1 (322)	
File pointers in Normal Play File arranged in descending order	File pointers in Trick Play File arranged in descending order
File pointers point to intra-coded data	File pointers point to inter-coded data

TABLE 2 (324)	
File pointers in Normal Play File arranged in descending order	File pointers in Trick Play File arranged in ascending order
File pointers point to intra-coded data	File pointers point to inter-coded data

5 When a subscriber is playing a file in a normal playback mode, data are spooled sequentially by alternating retrieval of data from the Normal Play File 300 and the Trick Mode File 320. When a trick play mode of fast forward is initiated, a location in the Trick Mode File 320 is identified by finding the closest file pointer corresponding to the current file pointer by reference to **TABLE 1**. Data are then spooled only from the trick
10 play file in the order dictated by the file pointers in **TABLE 2**.

 In a similar manner, when a subscriber is playing a file in a normal playback mode, data are spooled sequentially from both the Normal Play File 300 and the Trick Mode File 320. When a trick play mode of fast reverse is initiated, a location in the Trick Mode File 320 is identified by finding the closest file pointer to the current playback
15 point by reference to **TABLE 2**. Data are then spooled from the trick play file in the order dictated by the file pointers in **TABLE 2**.

 In either case, when the subscriber returns to normal play mode, the current file index in the Trick Mode File 320 is used as a starting location for normal play. Data are then pulled from both files 300 and 320 to produce normal playback. It is noted that
20 there is no overlap in the locations between the Normal Play file index and the Trick
Docket No.: SNY-T5712.02 PATENT

Mode indices. Playback will generally alternate between playing one or more I frames from file 320 and playing one or more B or P frames from file 300 to construct a complete set of the content.

As noted earlier, a similar result can be achieved with a single set of file indices such as that shown in **TABLE 3** (The trick play file pointers could be either ascending or descending.). In this example, fast forward trick play is achieved by playing out the trick play file 320 in the forward direction of the file pointers (top to bottom), and fast reverse trick play is achieved by playing out the trick play file 320 in the reverse direction of the file pointers (bottom to top). Again, normal playback involves selecting data from both files.

TABLE 3		FF	FR
File pointers in Normal Play File arranged in descending order	File pointers in Trick Play File arranged in descending order	↓	↑
File pointers point to intra-coded data	File pointers point to inter-coded data		

A process 330 for playback of content using the arrangement depicted in **FIGURE 8** is shown in **FIGURE 9** starting at 332. At 334, intra-coded frames are duplicated. At 336, "critical" intra-coded frames are selected for encryption. The inter-coded frames are stored in a first file (300) at 338. At 340, the selected frames are encrypted under the first CA encryption system. For multiple encryption systems, at 342 the duplicate frames are encrypted under a second CA encryption system. At 344 the intra-coded frames are stored in a second file and at 346, the duplicate intra-coded frames are stored in a third file. All or only part of the I frames may be encrypted within the second and third files. One or more index tables are created and stored that relate the first file to the second file and the first file to the third file at 348. In this example, a single

index table is depicted. When a subscriber initiates a playback at 350, a determination of playback mode is made at 352. If a normal playback mode has been invoked at 352, inter-coded frames from the first file and intra-coded frames from either the second or the third file (depending upon the decryption ability of the subscriber terminal 50) are
5 retrieved at 354 and assembled in forward sequence at 356 to produce full motion content. This content is then spooled to the output at 358 until the end is reached at 360 in which case the process stops at 362. If the end is not reached, control returns to 352 on a periodic or frequent basis to determine if the subscriber has invoked a trick mode.

If a trick mode has been invoked at 352, a location in the second file or third file
10 (depending upon the decryption ability of the subscriber terminal 50) is identified, by reference to **TABLE 3**, that is close to the current point of playback (e.g., the next inter-coded frame) at 364. If a fast forward trick mode has been invoked at 368, control passes to 372 where intra-coded frames are retrieved in forward order from the second or third file. If fast reverse trick mode has been invoked, control passes from 368 to 380 where
15 intra-coded frames are retrieved in reverse order from the second or third file. In either case, the retrieved frames are spooled to the output at 376. If the end of the file is reached at 388, the process stops at 362. Otherwise, control passes back to 352 to monitor the state of the selection of trick mode and either continue to operate in trick mode, change from one trick mode to the other or return to normal playback mode.

20 Many variations in this process are possible without departing from certain embodiments consistent with the present invention. For example, the ordering of certain actions can be rearranged without changing the basic operation. Also, equivalently, two tables such as **TABLE 1** and **TABLE 2** could be used. In this equivalent example, instead of designating an order of retrieval from the second or third file, the order is
25 always in the same direction, but with reference to a different table. Also in this variation, the tables used to determine entry points in the files for normal playback depends upon the trick mode selected, thus a mode determination is made to determine which table to use. Other variations will also occur to those skilled in the art upon consideration of the present teaching.

FIGURE 10 shows another embodiment of a playback process similar to that of **FIGURE 10**, but detailing certain variations. In this embodiment, processes preceding 352 are the same or similar to those of **FIGURE 9**. Also, to simplify the diagram, the end of file operation has been omitted, but adding it will be clear to those skilled in the art upon consideration of the present teaching.

In the normal play mode decision from 352, a determination is made as to whether or not the first (or next) frame for playback is located in the second or third file (depending upon the decryption capability of the subscriber terminal). If so, the next frame is retrieved from the second or third file at 408. If not, the next frame is retrieved from the first file at 410. In either event, the retrieved frame is spooled to the output at 412 and control returns to 344 to determine if a mode change has taken place. In other words, the presence or absence of an entry in the second or third file that corresponds to a next frame in the content is used to determine if content is retrieved from the first file at 410 on one hand or the second or third file on the other hand at 408.

When a fast forward trick mode is invoked at 368, intra-coded frames are retrieved from the second or third file (again depending upon the decryption ability of the subscriber terminal) in forward order at 420 and the frame is spooled to the output at 424. If no mode change occurs at 428, the process returns to 420 to retrieve the next frame. If the mode changes to normal playback mode at 428, control returns to 344.

If a fast reverse trick mode is invoked at 368, intra-coded frames are retrieved from the second or third file in reverse order at 440 and the frame is spooled to the output at 444. If no mode change occurs at 448, the process returns to 440 to retrieve the next frame. If the mode changes to normal playback mode at 448, control returns to 344.

If the mode changes to fast reverse at 428, control is passed to 440. If the mode changes to fast forward at 448, control passes to 420.

Again, many variations in this process are possible without departing from certain embodiments consistent with the present invention. For example, the ordering of certain actions can be rearranged without changing the basic operation, and end of file provisions should be provided. Also, equivalently, two tables such as **TABLE 1** and **TABLE 2**

could be used. In this equivalent example, instead of designating an order of retrieval from the second or third file at 420 and 440, the order is always in the same direction, but with reference to a different table. Also in this variation, the tables used to determine entry points in the files at 364 and for normal playback depends upon the trick mode selected, thus a mode determination is made at 364 to determine which table to use. Other variations including error trapping as well as other considerations will also occur to those skilled in the art upon consideration of the present teaching.

A preprocessor can be used to perform selective encryption of content to be loaded onto the VOD video server 22. A modified file protocol can be used to allow the VOD video server 22 to import and associate these files. Either the preprocessor or the VOD video server 22 can be used to perform the indexing. An alternate instantiation can be used to perform all selective encryption pre-processing within the video server itself. This can be accomplished by modifying the video server application to add a pre-processor task as a separate executable, called by the server during the process to prepare content for pre-encryption.

Additionally, in certain embodiments, this method overcomes the classic pre-encryption issue of supporting trick modes, but retains the other common problems of encryption "shelf life" and the additional handling required to prepare the stream for use on the VOD system.

Changes to the method employed by the VOD system for creating dynamic PSI data can be used to implement this architecture. The VOD system session manager 26 is made to be aware of which conditional access method is appropriate for a session requested by a specific subscriber in order to select the appropriate "critical" data file for the session. This information is in turn transferred to the VOD video server 22 that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The VOD video server 22 is cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in

addition to indicating the assigned PIDs for A/V, indicates the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

If one refers to the example movie scenario described above, the same movie
5 using 3.618GB of storage in the clear VOD state would require 3.159GBytes to store
using dynamic composition pre-encryption supporting two different CA systems – a
savings of almost 0.5 GB.

In accordance with certain embodiments consistent with the present invention,
certain of the functional blocks used to implement the VOD system can be implemented
10 using a programmed processor such as a general purpose computer. One example of
such a functional block is the session manager 26. However, the invention is not limited
to such exemplary embodiments, since other embodiments could be implemented using
hardware component equivalents such as special purpose hardware and/or dedicated
processors. Similarly, general purpose computers, microprocessor based computers,
15 micro-controllers, optical computers, analog computers, dedicated processors, application
specific circuits and/or dedicated hard wired logic may be used to construct alternative
equivalent embodiments.

Certain embodiments described herein, are or may be implemented using a
programmed processor executing programming instructions that are broadly described
20 above in flow chart form that can be stored on any suitable electronic or computer
readable storage medium and / or can be transmitted over any suitable electronic
communication medium. However, those skilled in the art will appreciate, upon
consideration of the present teaching, that the processes described above can be
implemented in any number of variations and in many suitable programming languages
25 without departing from embodiments of the present invention. For example, the order of
certain operations carried out can often be varied, additional operations can be added or
operations can be deleted without departing from certain embodiments of the invention.
Error trapping can be added and/or enhanced and variations can be made in user interface

and information presentation without departing from certain embodiments of the present invention. Such variations are contemplated and considered equivalent.

Those skilled in the art will appreciate, upon consideration of the above teachings, that the program operations and processes and associated data used to implement certain of the embodiments described above can be implemented using disc storage as well as other forms of storage such as for example Read Only Memory (ROM) devices, Random Access Memory (RAM) devices, network memory devices, optical storage elements, magnetic storage elements, magneto-optical storage elements, flash memory, core memory and/or other equivalent volatile and non-volatile storage technologies without departing from certain embodiments of the present invention. Such alternative storage devices should be considered equivalents.

Thus, in certain embodiments consistent with the present invention, a computer readable storage device for storing digital video content has at least one computer readable storage medium. A first file is stored on the storage medium containing unencrypted inter-coded frames of the digital video content. A second file is stored on the storage medium containing intra-coded frames of the digital video content encrypted under a first encryption algorithm. A third file is stored on the storage medium containing intra-coded frames of the digital video content encrypted under a second encryption algorithm. A first reference table relates frames in the first file to frames in the second file. A second reference table relates frames in the first file to frames in the third file.

While certain illustrative embodiments have been described, it is evident that many alternatives, modifications, permutations and variations will become apparent to those skilled in the art in light of the foregoing description.

What is claimed is: